Intelligence ≠ Autonomy ≠ Capability

J. P. Gunderson, L.F. Gunderson
Gamma Two, Inc.
{jgunders, lgunders}@gamma-two.com

ABSTRACT

In recent years, there has been a tendency to confuse the terms *Intelligence*, *Autonomy*, and *Capability*. In this paper we present the viewpoint that intelligence and capability are independent. These two factors describe an orthogonal design space that places upper bounds on the autonomy of the intelligent system. This design space for intelligent systems is illustrated by describing existing intelligent systems (some artificial, some natural) which demonstrate discrete points in the design space. Further, exemplars from biological systems indicate that there are natural constraints within the design space for intelligent systems, which suggest the need to balance the intelligence and capabilities of the designed systems. This design space is used to construct guidelines for the development of intelligent, capable, and autonomous systems.

KEYWORDS: *Intelligence, Autonomy, Capability, Intelligent Systems, Robotics, Autonomous Systems.*

1. INTRODUCTION

In recent years, there has been a tendency to confuse the terms *Intelligence*, *Autonomy*, and *Capability*. In this paper we present several arguments, drawn from research in biology, cognitive science, artificial intelligence, and psychology, that demonstrate that these concepts are far from equivalent. Further, we propose that, while these concepts are clearly related, they may be orthogonal. If this orthogonality can be demonstrated, then it could be leveraged to allow researchers to correctly categorize both the requirements of complex tasks, and the types of intelligent systems that are needed to achieve these tasks.

In this paper we present the viewpoint that intelligence and capability are independent – that one can have significant intelligence and lack capability, or vice versa. This is a critical design concept for engineers and scientists involved in the design and deployment of intelligent systems, since the design problem requires the careful application of the correct resources to solve specific problems. Autonomy is presented as an ability that is bounded above by the independent terms of intelligence and capability, however below this boundary, autonomy is shown to be conditionally independent as well.

This design space for intelligent systems is illustrated by describing existing intelligent systems (some artificial, some natural) which demonstrate discrete points in the design space. These exemplars are further used to extract design criteria for the needed levels of intelligence and capability, and the degree of autonomy which is necessary for an intelligent system to achieve its goals in its target environment.

2. DOMAIN

Before addressing the working definitions, it is necessary to develop a context. In this paper, the domain is considered to be goal directed behavior in dynamic and uncertain domains. Goal directed means that the intelligent systems deployed into these domains have goals that need to be achieved. This means that the systems are not behaving randomly, rather they are developing action sequences or selecting behaviors to change the state of the environment, and acting on these action sequences. The environment is dynamic, in that it is changing over time, in ways that are independent of the actions of the intelligent system. The intelligent system's perception of the environment is uncertain and the results of the actions taken by the intelligent system are non-deterministic.

These characteristics of the domain place significant demands on any intelligent system. In this domain, the intelligent system can only achieve probabilistic goal satisfaction; there is no guaranteed optimal performance. However, any system deployed in the real world must be able to deal with these problems.

3. INTELLIGENCE AND CAPABILITY

In order to have a common framework for discussion, it is necessary to have at least a working definition of the terms. While it would be nice to have precise definitions, these working definitions are meant to be tools for analyzing systems, not an absolute characterization that divides the world into classes.

In the area of measuring the performance of intelligent systems, a common approach is to measure the probability of goal satisfaction as a measure of intelligence. However, this results in confounding the effects of intelligence and capability. This confounding effect has contributed to the confusion of these two concepts, which has caused problems in the design and development of intelligent systems. The following sections present definitions for the terms intelligence and capability.

3.1 INTELLIGENCE

Intelligence has defied formal definition for as long as the concept has existed. There is a general consensus that intelligence is related to being able to solve problems, or to produce things that are of value to the society [1]. This

Report Documentation Page		*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE **AUG 2004**	2. REPORT TYPE	3. DATES COVERED **00-00-2004 to 00-00-2004**
4. TITLE AND SUBTITLE **Intelligence ≠ Autonomy ≠ Capability**		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Gamma Two, Inc,209 Kalamath St,Denver,CO,80223**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**

13. SUPPLEMENTARY NOTES **Proceedings of the 2004 Performance Metrics for Intelligent Systems Workshop (PerMIS ?04), Gaithersburg, MD on August 24-26 2004**

14. ABSTRACT **see report**

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**7**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

concept circles around the core idea that an intelligent system has the ability to achieve goals within an environment that is dynamic and uncertain. While intelligence is often characterized as a single entity, (e.g., Spearman's general intelligence, *g*), this theory has been criticized on a number of grounds [2]. Thomson argued in 1939 that there was no evidence that *g* represents any underlying structure in the nervous system of humans [3]. While Spearman proposed a two factor analysis, a combination of *g* and a collection of specific factors (*s*), recent work has proposed anywhere between seven and one hundred and fifty separate factors [4].

In the transition from natural intelligence to artificial intelligence, there is less emphasis on deriving the structural factors and more emphasis on the functional aspects. In part, this is due to the constructive nature of artificial intelligence as an engineered product, which is designed to meet specific requirements. Therefore, for this paper, the focus is on measuring intelligence via factor based performance metrics – how well does the intelligent system *develop* solutions to problems in a dynamic and uncertain environment. This is more closely allied with Newell and Simon's "intelligence as problem solving." [5] Albus and Meystel have proposed defining intelligence as "the ability to behave appropriately in an uncertain environment," where appropriate behavior will maximize the likelihood of goal satisfaction [6].

A critical aspect of intelligence is that it is based on the ability to determine or develop a good solution to a problem, not necessarily in the ability to execute it. In the human realm, a skilled mechanic, who can quickly and accurately diagnose a complex problem, does not become less intelligent because an injury prevents her from physically manipulating the wrench needed to execute the repair. Nor does she regain intelligence when the injury heals. Any working definition of intelligence must not fall into the trap where breaking one's leg makes one stupider. The definition used here is:

> **Intelligence**: the ability to determine behavior that will maximize the likelihood of goal satisfaction in a dynamic and uncertain environment.

This definition meets the criteria that an incidental change that impairs the ability of the system to execute the behaviors does not alter the intelligence of the system. Clearly however, the ability to successfully execute the appropriate behavior does affect the system's ability to satisfy goals in the real world. This successful execution is the capability of the system.

3.2 CAPABILITY

Execution capability has been less studied. Capability describes the ability of the intelligent system to successfully execute behaviors. An intelligent system may be able to correctly determine a valid course of action to achieve a goal, but be incapable of executing that course of action, while another system (such as a teleoperated robot) might be incapable of developing any 'intelligent' solution, however, given one, it can execute it and respond to minor failures during execution. An example of this would be our injured mechanic. She can correctly determine both the cause of a problem and the necessary repairs, but lack the capability to execute the repair. Her assistant, who may not have the intelligence or experience to solve the problem, can follow her instructions to effect the repair. The capability of her assistant does not imply intelligence.

> **Capability**: the ability to successfully execute behaviors or actions in a dynamic and uncertain environment.

There is one thing that should be noted about this definition. Unlike intelligence, capability is not goal oriented. This means that doing the wrong thing successfully does not imply reduced capability. However, selecting the wrong behavior to execute does imply reduced intelligence.

3.3 INDEPENDENCE OF INTELLIGENCE AND CAPABILITY

No attempt will be made in this paper to present a mathematical proof of independence. Rather, examples of intelligent systems which span the range of intelligence and capability will be presented. It is clear from the working definitions presented above, that intelligence is defined by the determination of behaviors, and has no definitional component related to execution. Likewise, capability is defined without any reference to the appropriateness, or correctness, of the actions, and is dependent only on the successful execution of the required actions.

Consider the example of vacuum cleaning robots and a two story home. The vacuum cleaning robots are available in several versions. They can be intelligent, with detailed maps of the furniture, and an understanding of traffic patterns; or they can be simple reactive systems that bump into things and vacuum at random. Second, they can be equipped with simple wheels, or they can be equipped to climb stairs as well. The goal is to keep all the carpets in the home clean.

Table 1 - Goal Satisfaction for vacuum cleaning robots

		Stairs	
		Low	High
Intelligence	Low	0.25	0.50
	High	0.50	1.0

In this case, the intelligent vacuum cleaners do a much better job of cleaning the corners, and getting the high-traffic areas, so the goal satisfaction is higher for the areas that they can reach. The stair climbing robots can reach more of the carpets, regardless of how good a job they do, so the goal satisfaction is higher overall. In effect, for all four of these robots goal satisfaction is a function of both intelligence and capability – but high intelligence can exist independently of high capability, and vice versa. This implies that:

$$g = f\,(c,i), \hspace{2cm} [1]$$
where   g = goal satisfaction,
        c = capability, and
        i = intelligence.

4. RELATIONSHIP OF AUTONOMY TO INTELLIGENCE AND CAPABILITY

In the previous section, working definitions for intelligence and capability were presented. These definitions offer several advantages in that they are independent, measurable, and are derived from the common usage of the words. As mentioned in the introduction, these are intended as tools to analyze systems, not necessarily to characterize them.

In this section autonomy is addressed, building from the previous definitions. The intent is to delineate the scope and relationships between these terms.

4.1 DEFINING AUTONOMY

The concept of autonomy, like that of intelligence, is controversial in the artificial intelligence and robotics communities. However, it is somewhat less controversial in everyday usage. According to one dictionary definition (in all cases the dictionary used is the 1969 American Heritage [7]). Autonomy is defined as:

Autonomy: 1.The condition or quality of being self-governing. 2. Self government, or the right of self-government; self-determination, independence. 3. A self-governing state, community, or group

The common term in all these definitions is *self-governing* or *self-government*. Clearly, to be autonomous is to be able to govern oneself, but what does this mean when applied to an intelligent system?

Govern also has an everyday meaning, which, while not commonly applied to a robot, has direct applicability to intelligent systems.

Govern: 1. To control the actions or behavior of; guide; direct. 2. To make and administer public policy for (a political unit); exercise sovereign authority in. 3. To control the speed of magnitude of; regulate. 4. To keep under control; restrain. 5. To decide; determine.

And the list goes on to grammatical uses. However, the clear thread in the definition of govern is the ability to decide and implement decisions. By extension, the sense of self-governing is the ability to decide and implement decisions for and by oneself. Hence, autonomy is the ability of a system to make choices and enforce its decisions. While this is consistent with the common usage, and has been supported by many researchers, this definition can be considered to conflict with a definition that limits the autonomy to be with respect to some goal or task assigned by an outside agency [8]. This definition addresses the issue raised by researchers such as Alan Schultz that a truly autonomous robot would be sitting on a beach somewhere, drinking motor oil, not slaving away on some human assigned task [9].

As an interesting side note, the etymology of govern is from the Greek *kubernan*: to steer, guide. This is the same root work that gives us cybernetics, coined by Norbert Weiner to refer to the theoretical study of control systems. In effect, an autonomous intelligent system is simply one that has the capability to control itself, to make decisions and implement those choices. Where, then, do these choices come from? Certainly, if the system is incapable of generating multiple options to achieve a goal, it cannot decide which behavior to undertake. With no viable behaviors, or a single solution to a problem, there is no choice – and therefore no ability to decide. While a system that can produce fifty possible goal satisfying solutions may be more intelligent that one that can only produce ten such solutions, it has no more autonomy. A system that can successfully execute more complex behaviors may be more capable than another, but that does not make it more autonomous, if an outside agency can override the chosen behavior and force an alternate behavior.

4.2 ESSENTIAL RELATIONSHIP

Notice there is nothing in the definition of autonomy that addresses the quality of the decisions made by the autonomous system. An entity can be autonomous and stupid, or be autonomous and smart. A system can make good decisions and implement them poorly, or implement them well. Autonomy is independent of intelligence (the ability to select appropriate behavior) and independent of capability (the ability to successfully execute an action or behavior.

Given the definition that autonomy is simply the ability of an entity to decide its own behavior and to execute that behavior, how does that relate to an intelligent system such as a robot? One critical factor is that the system must have the ability to select between options. If a system has no choice in the behaviors that it exhibits, then it cannot be

autonomous. In addition, if the choices made by the system can be overridden by an external agency, then the system is not autonomous. If a system has options to select from, and the ability to select and implement an option, then it is autonomous. Thus, a bi-metallic strip thermostat has no autonomy. The user selects the set point, and physics define the only possible action at any time. The thermostat is not free to say "Well, I know my set point is 72 degrees Fahrenheit, but I won't kick over until it gets down to 68." On the other hand, a simple random walk robot, with no goals except to keep moving, may have complete autonomy in the choice of direction and distance, and can autonomously fall down a staircase if it selects the wrong option.

Since the inability to produce a course of action precludes its use, and the inability to execute a course of action prevents it from being valid; it follows that intelligence and capability act as upper bounds on the autonomy of a system. However, while these two capacities limit the maximum autonomy of a system, below this limit the system can have a much or as little autonomy as the designer (in the case of an artificial intelligent system) allows.

Based on the definitions, it seems that intelligence and capabilities define an action space, and autonomy is bounded by this space (See Figure 1).
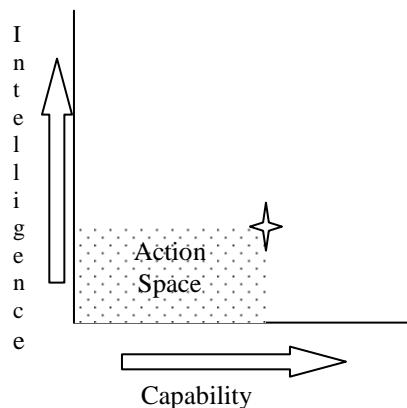

Figure 1: Effect of Intelligence and Capability on the range of autonomy. As intelligence and capability increase, the range of available options increase. The system can be autonomous, in which case it can select from these options. Or the system may not be autonomous, in which case its action space collapses to a single point for any goal/environment combination.

The available actions may be discrete options such as 'open a door' or 'send an email'. In the case of a mobile robot selecting a new heading or selecting a distance to travel the options may be drawn from an effectively continuous range. In terms of autonomy, it does not matter how many options are available, or even if there are an infinite number of options. All that matters is that there are options, as defined by the intelligence and the capabilities of the system, and that the system can select between them.

This leads naturally to the following question: Given a fixed system that can produce multiple feasible solutions to one problem, but can only produce a single solution to another problem in the same domain, is the system autonomous in the former case, and not autonomous in the latter?

Unfortunately, this leads to a typical white box versus black box problem. What if the system has multiple options, but always selects the same one? How does this differ from a system that only has that single behavior available to it? Is one autonomous and the other not – and how can one tell from the outside? This is clearly another variation on the 'strong' versus 'weak' question that has plagued artificial intelligence, cognitive science, and philosophy since their inceptions. The focus of this paper is not to answer such a complex conundrum, but to look at the implications of the relationship between autonomy and intelligence and capability.

Simply put, if the system cannot generate options for behavior, then the autonomy of the system is compromised; if the system cannot select between options, then the autonomy of the system is compromised; and if the system can select one of several options, but the execution of the option is controlled by an outside agency, then the autonomy is compromised. Thus, given that the system is enabled to choose between options in the method of achieving its goals, the intelligence and the capability of the system act as bounds on the autonomy of the system.

These two controlling factors, intelligence and capability, are independent in the abstract sense presented so far, and as such there are no constraints on either range. In theory, an intelligent system could be awesomely intelligent, yet totally incapable of achieving anything. Alternatively, a system could be a dumb as a box of rocks, yet have the capabilities of the most advanced robot ever imagined. These do not seem to be reasonable combinations. Examining of the existing intelligent systems, one does not see either extreme; rather the exemplars seem to be grouped with roughly balancing capabilities and the intelligence to use those capabilities. In the next section, examples of successful intelligent systems, deployed into harsh, dynamic, and uncertain domains are examined, drawn from biological systems.

5. CONTROLLING FACTORS ON INTELLIGENCE, AND CAPABILITY.

Brains are expensive. [10] This means that, for a species to compete effectively, the increased intelligence must covey an increased survival advantage. In the following discussion, the relationship of intelligence and capability in natural intelligences will be examined. In order to talk

about animal intelligence, some discussion is needed of the types of intelligence. Animal intelligence has been partitioned into two major classes 'menu-driven' intelligence and 'social' intelligence [11].

The 'menu' in menu-driven intelligence does not refer to pull-down option selection in a Graphical User Interface. Rather it is based in the concept that biological organisms must eat to survive, and therefore those that can find and acquire more varieties of food (their menu) more effectively have higher survival rates. Since finding food in a dangerous, harsh environment is extreme problem solving, it drives the development of intelligence.

Social intelligence is intelligence that is oriented towards communicating between individuals of the same species. The need for cooperation in hunting and defense require the ability to both develop team-based solutions to problems, and the associated skills of communication and maintaining group dynamics. This type of intelligence corresponds to Gardner's Lingustic and social intelligences.

These two types of intelligence will be treated separately. Since most organism can not describe their thought processes, all experimental evidence can show is the combination of intelligence and capacity. Starting with 'menu driven' intelligence, it has been shown that distribution of foods acts as a stimulus for mental development in primates [12]. In effect, the need to cognitively maintain and track the recognition patterns, locations, acquisition techniques and risks of additional types of food, increases the cognitive demand on the system, and requires allocation of resources to cognition, hence more intelligence is needed.

From this, it can be argued that in a highly competitive environment, any increase in intelligence that is not matched by an appropriate increase in capability could lead to the extinction of a species. The argument follows thusly. Let us consider two species in the same ecosystem. Both of these species are confronted with a new food source. One of the species uses its expensive brain power to construct a feeding strategy that it is capable of implementing. The other species constructs a feeding strategy that it is not capable of implementing. Clearly they have both exerted energy coming up with a solution, but only one has received a reward. In an evolutionary setting the species that came up with an implementable strategy has a major advantage. If the advantage is great enough, the losing species may face extinction. It is an interesting side note on the nature of biological systems, that the idea of an organism coming up with an unworkable strategy is almost unthinkable. This is the result of living in a harsh evolutionary system, where unsuccessful adaptations die out quickly.

'Social' intelligence is one that humans are most familiar with. Certainly for our species, it may be one of the most important of the many types of intelligences. However, even in this type of intelligence, capability plays

an important role. In order for this type of intelligence to be expressed, there must be both a sender with an idea and the capacity to transmit this idea, and a receiver with the capacity to receive the idea. By the same argument as above, if an organism spends energy developing a social intelligence, but it lacks the capacity to share, that organism will be out-competed by another organism that does not expend the energy to develop the social intelligence, or has the capacity to communicate the ideas.

The notion of a species having a capability without the corresponding intelligence to use that capacity is hard to imagine. Try to envision a bird, for example, perfectly capable of flying, but unable to 'think' of flying away when a predator attacks. Such a situation could, perhaps, arise but it would not exist for very long under evolutionary pressures. Either a ground based species would out compete the bird, since it would not need to maintain the expensive, but useless, flight equipment; or a subset of the birds would develop the intellect to use flight, and out compete their stupid brethren, Clearly there are birds which do not fly, but those that are capable of flying have the intellect to use the capability when needed.

For both of these classes of intelligence that have been established for animal minds, the intelligence and the capacity to use that intelligence must develop together. In the competitive natural environment, an imbalance between these factors would lead to competitive failure of the species.
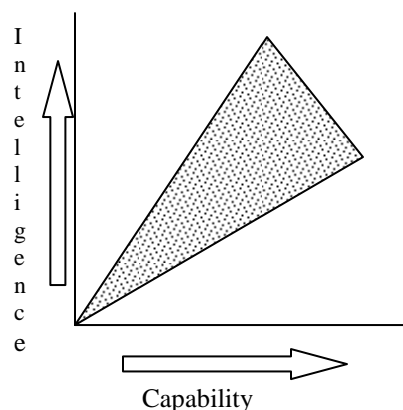

Figure 2: Apparent range of the balance between intelligence and capability in biological systems.

In Figure 2 is a rough representation of the existing biological intelligent systems, indicating the approximate parity between the capabilities of the system and the intelligence to use those capabilities. In general, the practical space is a close approximation to the abstract line where intelligence is precisely balanced by capability. However, there is a certain amount of variation as a result of the dynamic nature of biological systems. The region above the shaded area would correspond to species which had

intelligence in excess of their capabilities – resources invested in brains that could produce theoretic solutions to problems, solutions which could not be implemented with the existing capabilities. The region below the shaded area would correspond to species which had capabilities that it was impossible to figure out how to use.

In an environment with resource limitations and a high cost for failure the angle of the shaded region is fairly small. However, in the design space of engineered intelligent systems it is possible to create systems which are not bounded by the harsh realities of life. Researchers routinely develop intelligent systems that have capabilities which far outstrip their intelligence, and, less frequently, systems are designed and built that have the intelligence, but lack the capability to achieve their goals. In the following section, examples of these systems are presented and some approximate design rules are suggested to attempt to balance the allocation of resources between capability and intelligence.

5.1   LAW OF THE MINIMUM

In ecologic systems there is the notion that every species in an environment has a limiting factor. Leibig's Law of the Minimum states that for every species, there is a single limiting factor that controls growth of that species in that system [13]. For desert species, the limiting factor is the available water. For a rain forest species, the limiting factor is the amount of sun energy. This paper argues that, in an autonomous system, the limiting factor will either be intelligence or capability. If intelligence is the limiting factor, then increasing the capability will not significantly improve the performance of the system or allow for greater autonomy. Conversely, if capability is the limiting factor, then increasing the intelligence will not yield a significant improvement.

6.   GUIDELINES FOR DESIGN OF AUTONOMOUS SYSTEMS

The current state of the development of intelligent systems abounds with examples where the intelligence is out of balance with the capabilities of the system. In many cases these systems are built as research tools, and the focus is on exploring only one of the two aspects. However, there are also examples of systems which are deployed into dynamic and uncertain domains and which are intended to achieve specific goals.

The domain abounds with tele-operated systems where the capabilities of the system far outstrip the available intelligence. Systems which were intended as 'force multipliers' require three to five human operators. In part this is due to investing in capability rather than intelligence, and then falling back onto humans to do the 'hard part.'

6.1   BALANCE CAPABILITY AND INTELLIGENCE

This will sound trivial, but if the system is not meeting its performance goals, determine what the problem is before attempting a fix. The reasoning above demonstrates that there are at least two possible, independent causes for reduced performance metrics. Either could be the limiting factor that is restricting performance. Fixing the wrong one won't help as much as fixing the right one.

Determine the limiting factor and address it, do not just throw more resources at the easy factor. In dynamic domains, where the intelligent system must deal with uncertainty, it is critical to limit the intelligence to that which is sufficient to solve the problem. Adding more capability to an intelligence limited system will show an improvement, but adding more intelligence will show a greater improvement.

6.2   DON'T MAKE IT AUTONOMOUS, UNLESS THAT IS WHAT YOU NEED.

Autonomy is problematic. If the system should be tightly controlled, then it needs to be tightly controlled. Autonomy is not a 'magic bullet' that will make an ineffective system work in a dynamic, uncertain domain. If the problem domain requires a specific response to specific inputs, then the designer must provide that mapping. Rather than adding autonomy, focus on building the correct mapping.

If the system needs to be autonomous, and in dynamic, uncertain domains most successful systems must be autonomous[14], the intelligence and capability must be sufficient to support the autonomy. It is pointless to give the system autonomy if it has neither the brains nor the capability to use the autonomy successfully. In addition, making the system autonomous will not magically enable it to do its job. Autonomous stupidity is easy.

7.   CONCLUSIONS

One drawback to performance based metrics of intelligence is the confounding effects of intelligence and capability. No one would argue that a physicist such as Steven Hawking is not intelligent. Yet if he were put into a room and given a battery of intelligence tests, his physical disabilities would impact the measure of his intelligence – unless the tests were specifically designed to correct for the confounding effects.

Recently, it has been the practice to measure the performance of intelligent systems and label the result 'intelligence'; however, the tests measure the combination of intelligence and capability. This has resulted in the merging of these two very important, very different terms. It is not uncommon to hear researchers say that capability is the same as intelligence. This paper is an attempt to clarify these two terms, so that they can be used to more effectively describe the complexities of intelligent systems.

Intelligence is a cognitive process that allows a system to propose a viable solution to a problem or task. Capability is the ability to implement or execute a proposed solution in a dynamic, uncertain environment successfully. Both of these abilities are necessary for goal satisfaction and both are very difficult. However, by conflating the two terms, researchers run the risk of having a system that fails to achieve its goals due to insufficient intelligence, but increasing the capacity of the system in an attempt to improve it. This can be both fruitless and frustrating for the researcher.

By keeping the concepts of capability and intelligence orthogonal, the system designer has a design tool that will allow her to determine the correct area to focus attention, and the correct type of improvements that will result in a better system.

8. REFERENCES

[1] Gardner, H. *Frames of Mind: The Theory of Multiple Intelligences Tenth anniversary edition*, New York: Basic Books, 1993, ch.4, pp. 60-62.

[2] Richardson, K, *The Making of Intelligence*, New York: Columbia University Press, 2000, ch.1, pp. 8-10.

[3] Thomson, G.H. *The Factorial Analysis of Human Ability*, Boston: Houghton Mifflin, 1939, ch.1, pp.4.

[4] Gardner, H., Kornhaber, M.L., and Wake, W. K. *Intelligence: Multiple Perspectives*, Fort Worth, TX: Harcourt Brace College Publishers, 1996, ch.3, pp. 70-75.

[5] Newell, A. and Simon, H. A. *Human Problem Solving* Englewood cliffs, New Jersey: Prentice –Hall, Inc., 1972, ch. 1, pp. 6–9.

[6] Albus, J. S. and Meystel, A. M. *Engineering of Mind*, NewYork: Wiley Interscience, 2001, ch.1, pp. 6-7.

[7] *American Heritage Dictionary of the English Language*, New York: American Heritage Publishing Co., 1969.

[8] Barber, K.S. and Martin. C.E., "Agent Autonomy: Specification, Measurement, and Dynamic Adjustment," in Proceedings of the Autonomy Control Software Workshop, part of Autonomous Agents, 1999, pp.8-15.

[9] Schultz, Alan, Panel Discussion, AAAI Spring Symposium on Interaction Between Humans and Autonomous Systems over Extended Operation, 2004.

[10] Gunderson, J. P. and Gunderson, L. F. "Mom! The Vacuum Cleaner is Chasing the Dog Again," Proceedings of the 2003 PerMIS Workshop. 2003.

[11] Yoerg, S. I., *Clever as a Fox*, New York: Bloomsbury, 2001, ch. 8, pp. 162-175

[12] Milton, K. "Distribution Patterns of Tropical Plant Food as a Stimulus tp Primate Mental Development," *American Anthropologist*, vol. 83, pp. 534-548.

[13] Barbour, M.G., Burk, J.H., and Pitts, W.D. *Terrestrial Plant Ecology*, Menlo Park: CA, 1987, ch.3, pp.29-30.

[14] Gunderson, J.P. and Martin, W. N. "A Probaility-Aware Planning Architecture to Support Variable Autonomy," *American Association for Artificail Intelligence Technical Report SS-03-04*, Menlo Park, CA: AAAI Press, pp. 83 – 88.